

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 05-12-2017		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Aug-2014 - 31-Jul-2017	
4. TITLE AND SUBTITLE Final Report: A Belief-Space Approach to Integrated Intelligence- Research Area 10.3: Intelligent Networks			5a. CONTRACT NUMBER W911NF-14-1-0433		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Massachusetts Institute of Technology (MIT) 77 Massachusetts Avenue NE18-901 Cambridge, MA 02139 -4307			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 65633-NS.13		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Leslie Kaelbling
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 617-258-9695

RPPR Final Report

as of 23-Jan-2018

Agency Code:

Proposal Number: 65633NS

Agreement Number: W911NF-14-1-0433

INVESTIGATOR(S):

Name: Leslie P. Kaelbling
Email: lpk@csail.mit.edu
Phone Number: 6172589695
Principal: Y

Organization: **Massachusetts Institute of Technology (MIT)**

Address: 77 Massachusetts Avenue, Cambridge, MA 021394307

Country: USA

DUNS Number: 001425594

EIN: 042103594

Report Date: 31-Oct-2017

Date Received: 05-Dec-2017

Final Report for Period Beginning 01-Aug-2014 and Ending 31-Jul-2017

Title: A Belief-Space Approach to Integrated Intelligence- Research Area 10.3: Intelligent Networks

Begin Performance Period: 01-Aug-2014

End Performance Period: 31-Jul-2017

Report Term: 0-Other

Submitted By: Tomas Lozano-Perez

Email: tlp@mit.edu

Phone: (617) 253-7889

Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 1

STEM Participants: 6

Major Goals: The overall goal of this work is to develop the perception, estimation, planning, and control techniques necessary to enable autonomous agents to perform robustly and intelligently in complex uncertain domains. This includes the ability to intelligently interact and coordinate with humans or other agents so as to achieve goals effectively and efficiently.

Specifically our goal has been to develop algorithms for effective and efficient planning in domains that are characterized by a variety of action types, each with continuous parameters, for example the variety of manipulation actions available to a robot (picking, placing, pushing, tilting, etc.). Furthermore, the actions take place in the presence of uncertainty both as to the current state of the world and as to the actual result of the action.

Accomplishments: A summary of our most important results, embodied in our conference publications is included in the attached PDF file. Copies of the full papers have been uploaded.

Training Opportunities: Throughout the period of this award students, both graduate and undergraduate, played a key role in the research. The students

- o met with the PIs for an hour every week
- o participated in a weekly meeting in which students presented progress and received feedback from the research group.
- o wrote papers on their research and submitted them to leading conferences

Results Dissemination: The research results have been published in leading conferences and journals. The papers, including code, has been made available on our group's public web site.

The PIs have given talks at conferences and invited talks at leading research universities (CMU, U. Washington, UIUC, etc).

The PIs have done presentations aimed at minority students at MIT and also presentations to High School students and also High School Teachers.

Honors and Awards: Nothing to Report

RPPR Final Report
as of 23-Jan-2018

Protocol Activity Status:

Technology Transfer: Nothing to Report

PARTICIPANTS:

Participant Type: PD/PI

Participant: Leslie Kaelbling

Person Months Worked: 3.00

Funding Support:

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Participant Type: Co PD/PI

Participant: Tomas Lozano-Perez

Person Months Worked: 3.00

Funding Support:

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Participant Type: Graduate Student (research assistant)

Participant: Caelan Garrett

Person Months Worked: 6.00

Funding Support:

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Participant Type: Graduate Student (research assistant)

Participant: Ashwin Deshpande

Person Months Worked: 12.00

Funding Support:

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Participant Type: Graduate Student (research assistant)

Participant: Ariel Anders

Person Months Worked: 12.00

Funding Support:

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Participant Type: Graduate Student (research assistant)

Participant: Zi Wang

Person Months Worked: 12.00

Funding Support:

RPPR Final Report

as of 23-Jan-2018

Project Contribution:
International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

Participant Type: Undergraduate Student

Participant: Brian Axelrod

Person Months Worked: 3.00

Funding Support:

Project Contribution:
International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

Participant Type: Graduate Student (research assistant)

Participant: Kenji Kawaguchi

Person Months Worked: 6.00

Funding Support:

Project Contribution:
International Collaboration:
International Travel:
National Academy Member: N
Other Collaborators:

ARTICLES:

Publication Type: Journal Article

Peer Reviewed: Y

Publication Status: 2-Awaiting Publication

Journal: The International Journal of Robotics Research

Publication Identifier Type: DOI

Publication Identifier: 10.1177/0278364917739114

Volume:

Issue:

First Page #: 027836491

Date Submitted: 12/4/17 12:00AM

Date Published: 11/1/17 4:00AM

Publication Location:

Article Title: FFRob: Leveraging symbolic planning for efficient task and motion planning

Authors: Caelan Reed Garrett, Tomás Lozano-Pérez, Leslie Pack Kaelbling

Keywords: Task and motion planning, manipulation planning, AI reasoning

Abstract: Mobile manipulation problems involving many objects are challenging to solve due to the high dimensionality and multi-modality of their hybrid configuration spaces. Planners that perform a purely geometric search are prohibitively slow for solving these problems because they are unable to factor the configuration space. Symbolic task planners can efficiently construct plans involving many variables but cannot represent the geometric and kinematic constraints required in manipulation. We present the FFRob algorithm for solving task and motion planning problems. First, we introduce Extended Action Specification (EAS) as a general purpose planning representation that supports arbitrary predicates as conditions. We adapt existing heuristic search ideas for solving \proc{strips} planning problems, particularly delete-relaxations, to solve EAS problem instances. We then apply the EAS representation and planners to manipulation problems resulting in FFRob. FFRob iteratively discretizes task a

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

CONFERENCE PAPERS:

RPPR Final Report
as of 23-Jan-2018

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: International Joint Conference on Artificial Intelligence
Date Received: 04-Dec-2017 Conference Date: 25-Jul-2015 Date Published: 25-Jul-2017
Conference Location: Buenos Aires, Argentina
Paper Title: Symbol Acquisition for Probabilistic High-Level Planning
Authors: George Konidaris, Leslie Kaelblin, Tomas Lozano-Perez
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: IEEE/RSJ International Conference on Intelligent Robots and Systems
Date Received: 04-Dec-2017 Conference Date: 28-Sep-2015 Date Published: 28-Sep-2015
Conference Location: Hamburg, Germany
Paper Title: Backward-Forward Search for Manipulation Planning
Authors: Caelan Garrett, Tomas Lozano-Perez, Leslie Kaelbling
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: IEEE/RSJ International Conference on Intelligent Robots and Systems
Date Received: 04-Dec-2017 Conference Date: 28-Sep-2015 Date Published: 28-Sep-2015
Conference Location: Hamburg, January
Paper Title: Hierarchical planning for multi-contact non-prehensile manipulation
Authors: Gilwoo Lee, Tomas Lozano-Perez, Leslie Kaelbling
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: Int. Joint Conf. on Artificial Intelligence
Date Received: 28-Jul-2016 Conference Date: 09-Jul-2016 Date Published: 09-Jul-2016
Conference Location: New York City
Paper Title: Learning to rank for synthesizing planning heuristics
Authors: Caelan Reed Garrett, Leslie Pack Kaelbling, Tomas Lozano-Perez.
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: IEEE Conference on Robotics and Automation
Date Received: 28-Jul-2016 Conference Date: 16-May-2016 Date Published: 16-May-2016
Conference Location: Stockholm, Sweden
Paper Title: Implicit belief-space pre-images for hierarchical planning and execution
Authors: Leslie Pack Kaelbling, Tomas Lozano-Perez
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: IEEE Conference on Robotics and Automation
Date Received: 28-Jul-2016 Conference Date: 16-May-2016 Date Published: 16-May-2016
Conference Location: Stockholm, Sweden
Paper Title: Searching for physical objects in partially known environments
Authors: Xinkun Nie, Lawson L.S. Wong, Leslie Pack Kaelbling
Acknowledged Federal Support: **Y**

RPPR Final Report
as of 23-Jan-2018

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: International Joint Conference on Artificial Intelligence
Date Received: 28-Jul-2016 Conference Date: 09-Jul-2016 Date Published: 09-Jul-2016
Conference Location: New York City
Paper Title: Object-based world modeling in semi-static environments with dependent Dirichlet process mixtures
Authors: Lawson L.S. Wong, Thanard Kurutach, Tomas Lozano-Perez, Leslie Pack Kaelbling.
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: Robotics: Science and Systems
Date Received: 04-Dec-2017 Conference Date: 12-Jul-2017 Date Published: 12-Jul-2017
Conference Location: Boston, Mass
Paper Title: Sample-based methods for factored task and motion planning
Authors: Caelan Garrett, Tomas Lozano-Perez, Leslie Kaelbling
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: International Conference on Artificial Intelligence and Statistics
Date Received: 04-Dec-2017 Conference Date: 19-May-2016 Date Published: 19-May-2016
Conference Location: Cadiz, Spain
Paper Title: Optimization as Estimation with Gaussian Processes in Bandit Settings
Authors: Zi Wang, Bolei Zhou, Stefanie Jegelka
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: AAAI Conference on Artificial Intelligence
Date Received: 04-Dec-2017 Conference Date: 12-Feb-2016 Date Published: 12-Feb-2016
Conference Location: Phoenix, Arizona
Paper Title: Bounded optimal exploration in MDP
Authors: Kenji Kawaguchi
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: Robotics: Science and Systems
Date Received: 05-Dec-2017 Conference Date: 26-Jun-2017 Date Published: 26-Jun-2017
Conference Location: Boston, Mass
Paper Title: Provably safe robot navigation with obstacle uncertainty.
Authors: Brian Axelrod, Leslie Kaelbling, Tomas Lozano-Perez
Acknowledged Federal Support: **Y**

Interim Project Report for Contract W911NF1410433

Report Period - Start Date: 08-01-2014; End Date: 07-31-2017

Statement of the problem studied

The overall goal of this work is to develop the perception, estimation, planning, and control techniques necessary to enable autonomous agents to perform robustly and intelligently in complex uncertain domains. This includes the ability to intelligently interact and coordinate with humans or other agents so as to achieve goals effectively and efficiently.

The focus of our work has been to develop algorithms for effective and efficient planning in domains that are characterized by a variety of action types, each with continuous parameters, for example the variety of manipulation actions available to a robot (picking, placing, pushing, tilting, etc.). Furthermore, the actions take place in the presence of uncertainty both as to the current state of the world and as to the actual result of the action.

In this report, we highlight results that have appeared in major conferences. Some of this work is on its way to appearing in journal publications as well.

Summary of the most important results

Backward-Forward Search for Manipulation Planning [4]

We are interested in solving manipulation planning problems in high-dimensional hybrid configuration spaces. A state of such a system is characterized by a finite set of configuration variables that may be discrete (such as which object a robot is holding or whether the light is turned on) or continuous (such as the joint-space configuration of a robot or the pose of an object).

Without making any assumptions about the nature of the configuration space and the transition dynamics, planning in such a space is quite difficult. We have developed a problem representation that can reveal useful underlying structure in the domain that will be exploited by our method. There are three important kinds of leverage:

- *Factoring and sparsity*: by representing the state space as the product of the spaces of a set of state variables, we are able to assert that each action of the robot affects only a small subset of the state variables, allowing individual actions to be contemplated in state spaces that are effectively much smaller.
- *Continuous modes*: there are some continuous subspaces of the whole space that have continuous dynamics, which allows us to use classic sample-based robot motion planning techniques to move within those subspaces.
- *Heuristic estimates*: by constructing relaxed versions of a planning problem, we can efficiently obtain estimates of the cost to reach a goal state and use these estimates to make the search for a solution much more efficient.

We have developed a new planning algorithm, HBF, and applied it to a variety of different manipulation problems (shown below) to characterize its performance. Solving these problems requires stacking, regrasping, pushing, and long-horizon manipulation. The planner and PR2 robot manipulation simulations were written in Python using OpenRAVE. In each problem, red objects represent moveable objects that have no particular goal condition. However, they impose geometric constraints on the problem and must, in many cases, be manipulated in order to produce a satisfying plan.




Figure 1: An early state on a valid plan for each problem.

Problem 1

The goal constraint is for the green block, which is surrounded by 8 red blocks, to be on the green region on the table. Notice that the right table has 40 movable red objects on it that do not block the path of the green object.

Problem 2

The goal constraint is for the green cylinder to be at the green point on the edge of the table. The cylinder is too big for the robot to grasp, so it must push it instead. The robot must move several of the red objects and then push the green cylinder several times to solve the problem.

Problem 3

A thin but wide green block starts behind a similar blue block. The goal constraints are that the green block be at the green point and that blue block be at the blue point, which is, in fact, its initial location. This problem is non-monotonic in that the robot must first violate one of the goal constraints, and then re-achieve it. Additionally, the right cubby housing the green goal point is thinner than the left cubby, so the green block can only be placed using a subset of its grasps, all of which are infeasible for picking it up at its initial location. This forces the robot to place and regrasp the green block.

Problem 4

The goal constraints are that the green block be on the green region of the table, the blue block be on the blue region of the table, and the black block be on top of the blue block. Because the black block must end on the blue block, which itself must be moved, no static pre-sampling of object poses would suffice to solve the problem. Additionally, a red block starts on top of the green block, preventing immediate movement of the green block.

Problem 5

This is exactly the same problem considered by Srivastava et al. in earlier work. The goal constraint is to be holding the red cylinder with an arbitrary grasp. 39 blue cylinders crowd the table, blocking the red object.

Problem 6

The goal constraints are that all 7 blue blocks must be on the left table and all 7 green blocks must be on the right table. There are also 14 red blocks. The close proximity of the blocks forces the planner to carefully order its operations as well as to move red blocks out of the way.

The experiments show that by leveraging the factored nature of common manipulation actions, HBF is able to efficiently solve complex manipulation tasks. The runtimes are improvements over runtimes on comparable problems reported in previous work. Additionally, the dynamic search allows HBF to solve regrasping, pushing, and stacking problems all using the same planning algorithm.

Hierarchical planning for multi-contact non-prehensile manipulation [9]

We have explored a hierarchical approach to planning sequences of non-prehensile and prehensile actions. Our planner operates hierarchically, first finding a sequence of qualitative “object contact states” that characterize which parts of the moving object are in contact with which parts of other objects, then finding a feasible sequence of poses for the object (figure 2), and finally finding a sequence of contact points for the manipulators on the object (figure 3). This hierarchical structure provides significant search guidance, and divides the problem into three search problems that are much smaller than a search in the full combined configuration space of the object and manipulators.




Figure 2: A contact state graph with poses connected through linear interpolation. Poses connecting two contact states are very close to each other.

To find a robot-contact plan, we discretize the object’s surface into a set of possible contact points and define a state to contain an object pose and a set of contacts of the robot’s manipulators on the object. We then identify states that are *feasible*: both *accessible*, meaning that the robot can reach all of the specified contacts and *stabilizable*, meaning that there exists a set of contact forces between the object and the robot’s manipulators, as well as the fixed objects, that can stabilize the object against gravity (figure 3).

Figure 4 illustrates the connected search spaces: within the discrete contact states in the contact-state graph, there are individual object poses, and a path through object-contact space can be realized by a path through object pose space. Then, for each object pose, there is a set of robot contacts, and a path through object pose space can be realized by a path of transit and transfer motions through the combined space of robot contacts and object poses.

We have implemented a version of this planner (in simulation) for planar objects and two robot contacts, without any further kinematic or collision constraints introduced to model the robot performing the manipulation. We tested these approaches on two problems. The first, shown in figure 5, focuses on sequencing non-prehensile manipulation steps. There is an obstacle in the middle of the table, and the goal in this problem is to move the box to the other side of the table. Allowing only nonprehensile manipulation, the planner is able to find a solution.

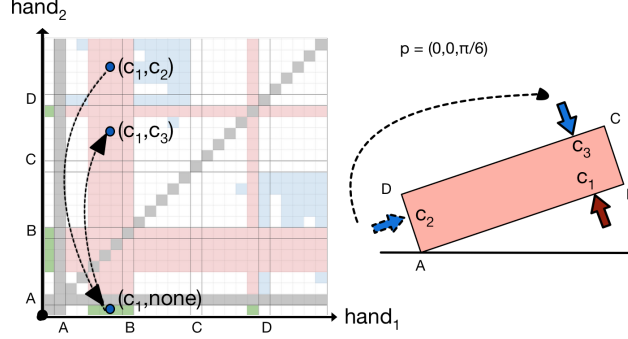


Figure 3: Robot contact space for $p = (0, 0, \pi/6)$. Each axis represents possible contact points along the object’s surface accessible by $hand_1$ and $hand_2$. The leftmost column and the bottom row represent no-contact for $hand_1$ and $hand_2$, respectively. Green cells represent feasible states with only one contact, i.e. where the object can be balanced by only one hand. If either hand makes the object stabilizable on its own, the other hand can place itself on any accessible surface; these states are colored in red. For example, if a row’s leftmost cell is green, all accessible cells in the row becomes red. States that require both hands are colored in blue. Grey cells represent invalid or inaccessible states. Since vertex A is already in contact with ground, any state containing A is inaccessible. White cells are infeasible. A *transit* is a transition from a red state to another red state in the same row or column. The example shows transits from (c_1, c_2) to $(c_1, none)$ to (c_1, c_3) , changing which manipulator is stabilizing the object.




Figure 4: The relationship between the spaces of object contacts, object poses, and robot contacts.

Symbol Acquisition for Probabilistic High-Level Planning [8]

Systems that combine high-level planning with low-level control are capable of generating complex, goal-driven behavior. But, they are hard to design because they require a difficult integration of symbolic reasoning and low-level motor control.

Recently, we showed how to automatically construct a symbolic representation suitable for planning in a high-dimensional, continuous domain. This work modeled the low-level domain as a semi-Markov decision process and formalized a propositional symbol as the name given to a grounding set of low-level states (represented compactly using a learned classifier). Their key result was that the symbols required to determine the feasibility of a plan are directly determined by characteristics of the actions available to an agent. This close relationship removes the need to hand-design symbolic representations of the world and enables an agent to, in principle, acquire them autonomously.

However, a set-based symbol formulation cannot deal with learned sets that may not be exactly correct, and can only determine whether or not the probability of successfully executing a plan is 1. These restrictions are ill-suited to the real-world, where learning necessarily results in uncertainty and all plans have some probability of failure




Figure 5: Key frames from a sample solution trajectory for tumbling one box over another. Red lines indicate the force direction of the robot contact pushing the object. The “hands” simply highlight the location of the chosen robot contacts.

In our new work, we introduced a probabilistic reformulation of symbolic representations capable of naturally dealing with uncertain representations and probabilistic plans. This is achieved by moving from sets and logical operations to probability distributions and probabilistic operations. We use this framework to design an agent that autonomously learns a completely symbolic representation of a computer game domain, enabling very fast planning using an off-the-shelf probabilistic planner.

Implicit Belief-Space Pre-images for Hierarchical Planning and Execution[6]

We have developed a method for planning and execution in very high-dimensional mixed discrete and continuous spaces in the presence of uncertainty, based on an implicit, factored approximation of pre-images in continuous spaces and have extended it to apply to the case where the space is actually the belief space of probability distributions over underlying world states. We have implemented a planning algorithm that searches in the space of pre-images in this representation. Finally, we have demonstrated this approach in a mobile-manipulation domain that combines pushing with pick-and-place manipulation using actions with motion and sensing error (see Figure 6).




Figure 6: A test case in which the object must be moved out of the way and the red object pushed to the left corner of the table. Initial situation is on the left; a final situation from a sample run is on the right.

This approach is related to our earlier work on the Hierarchical Planning in the Now (HPN) system. The two key differences are: (1) the representation of pre-images by using implicit fluents such as *CanReachHome* and *CanPlace* instead of using explicit representations of swept volumes of particular paths as obstacles, and (2) the introduction of a general notion of conditioning in the regression algorithm to handle these implicit fluents. Together these extensions generalize and make systematic the pre-image computation approach

in our earlier work. Our contribution is to show how to tractably plan using implicit representations of pre-images in hybrid state-spaces and to demonstrate their use for hierarchical planning and execution monitoring in real robot manipulation problems.

We have shown empirically that execution monitoring using pre-images provides a substantial improvement in computational efficiency over continual replanning, and that the hierarchical algorithm enabled by abstract pre-image backchaining provides even further efficiency improvements.

Searching for Physical Objects in Partially Known Environments[10]

As the perception, locomotion, and manipulation abilities of robots begin to improve, we begin to contemplate constructing robots that can help with household chores or assist in disaster recovery. In open domains such as these, robots will have to be able to operate in cluttered domains and be able to locate objects of interest within them.

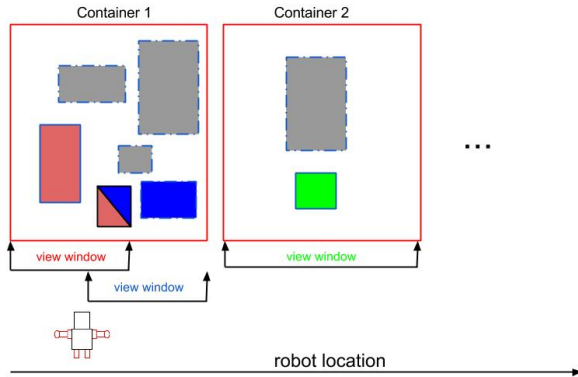


Figure 7: A robot searching for objects in a 2-D domain. The colored rectangles are the objects that can be observed by the robot if the robot sits within the corresponding view window. The half-blue half-red object can be observed both from the red view window and the blue view window. The grey shaded objects are occluded, and are only visible after objects in the front have been removed. In our problem, the robot is asked to obtain an object of a specific type, and it must choose both where to look and possibly what to remove in order to find the object. When the target object type is occluded, contextual cues such as collocated object types and spatial constraints can be used in an inference process to guide the search.

In this paper, we have addressed the problem of a mobile manipulation robot searching for an object in a cluttered domain that is populated with an unknown number of objects in an unknown arrangement (see Figure 7). The robot must move around its environment, looking in containers, moving occluding objects to improve its view, and reasoning about collocation of objects of different types, all in service of finding a desired object. We did not address issues of low-level perception or of manipulation in this paper; rather, we provide a general framework for reasoning about arrangements of unknown objects and for planning how to search effectively for a desired object.

The key contribution in reasoning is a Markov-chain Monte Carlo (MCMC) method for drawing samples of the arrangements of objects in an occluded container, conditioned on previously gathered information of a variety of types. Relevant information includes: spatial knowledge, such as the sizes and shapes of containers such as shelves, historical knowledge, of which objects have already been removed from the containers, type co-occurrence knowledge, which says which types of objects are most likely to occur near one another, and other global constraints in the domain.

The key contribution in planning is a receding-horizon forward search in the space of beliefs about (distributions over) arrangements (including number and type) of objects in the domain. In our domain, observations are drawn from a continuous space of object poses, so making this search tractable requires the construction of abstract observation models, which reduce the effective branching factor and number of state samples needed to represent beliefs.

Object-based World Modeling in Semi-Static Environments with Dependent Dirichlet Process Mixtures[12]

Robots need to know about objects in order to perform most tasks in human-centered environments. Objects should be understood in terms of semantic attributes such as type, pose, function, and possibly relations with other objects. Semantic perception tools are increasingly becoming available, and it is tempting to use them as black-box perception modules. However, such perception is still error-prone, due to noise, occlusion, clutter, and limited fields of view. To achieve greater reliability, our strategy is to aggregate the output from noisy perception pipelines, across time and space (different viewpoints), and estimate the true state, i.e., the world model (see Figure 8).

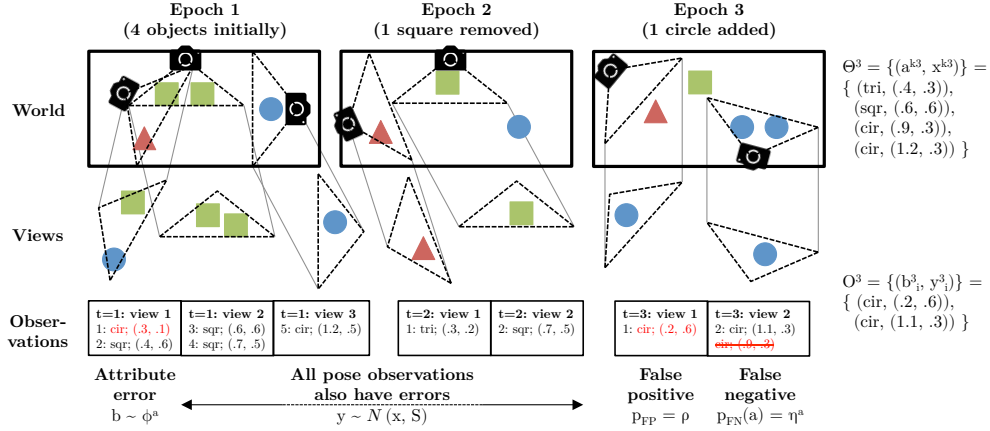


Figure 8: An illustration of the world modeling problem. An unknown number of objects exist in the world (top row), and change in pose and number over time (world at each epoch enclosed in box). At each epoch, limited views of the world are captured, as depicted by the triangular viewcones. Within these viewcones, objects and their attributes are detected using black-box perception modules (e.g., off-the-shelf object detectors). In this example, the attributes are shape type (discrete) and 2-D location. The observations are noisy, as depicted by the perturbed versions of viewcones in the middle row. Uncertainty exists both in the attribute values and the existence of objects, as detections may include false positives and negatives (e.g., $t = 3$). The actual attribute detection values obtained from the views are shown in the bottom row (Observations); this is the format of input data. Given these noisy measurements as input, the goal is to determine which objects were in existence at each epoch, their attribute values (e.g., Θ^3 in top right), and their progression over time.

Estimating properties of individuals from noisy observations is a relatively simple statistical estimation problem if the observations are labeled according to which individual generated them. Even when the underlying attributes of the individual change over time, estimating their history reduces to inference in a hidden Markov model.

The key difficulty is data association. We do not know which particular individual is responsible for each observation; determining an appropriate association of observations to individuals is key. The only information we have to make such associations are noisy and partial observations, which may contain errors both in attribute values and in number.

This problem was first addressed in the context of multiple-target tracking. A classical solution is multiple hypothesis tracking, which has been applied in previous world modeling applications. Others have pointed out drawbacks in using the MHT, which include inefficiency due to considering an exponential number of hypotheses, and the inability to revisit associations from previously-considered views (the MHT is a filtering algorithm). Inspired by this, others have proposed different Markov-chain Monte Carlo (MCMC) methods for data association, and have demonstrated superior tracking performance.

In multiple-target tracking problems, each targets state (typically location) changes between observations.

However, if we consider applications such as tracking objects in a household, the dynamics are different: most objects tend to stay in the same state when they are not being actively used. In this paper, we studied the world modeling problem in semi-static environments, where time is divided into known epochs, and within each epoch the world is stationary. Intuitively, data association should be easier within static periods, since there is no uncertainty arising from stochastic dynamics.

At the other end of the spectrum, in previous work we considered the world modeling problem under a static world assumption. We proposed a clustering-based view of the problem, where objects are treated as cluster components (in a joint attribute space), and observations are noisy measurements generated from these clusters. We used Bayesian nonparametric models to handle an unknown number of objects, in particular the Dirichlet process mixture model (DPMM). This approach is fundamentally limited by the DPMMs inability to capture temporal dynamics.

Dependent Dirichlet processes (DDP), in contrast, are capable of modeling dynamic clusters. We use a DDP mixture model to infer object attributes and their changes over time, including the addition and removal of objects in the world. A novel approximate MAP inference method is also proposed.

Learning to Rank for Synthesizing Planning Heuristics[3]

Forward state-space greedy heuristic search is a powerful technique that can solve large planning problems. However, its success is strongly dependent on the quality of its heuristic. Many domain-independent heuristics estimate the distance to the goal by quickly solving easier, approximated planning problems. While domain-independent heuristics have enabled planners to solve a much larger class of problems, there is a large amount of room to improve their estimates. In particular, the effectiveness of many domain-independent heuristics varies across domains, with poor performance occurring when the approximations in the heuristic discard a large amount of information about the problem.

Previous work has attempted to overcome the limitations of these approximations by learning a domain-specific heuristic correction. Yoon et al. formulated learning a correction for the FastForward (FF) heuristic as a regression problem and solved it using ordinary least-squares regression. While the resulting planner is no longer domain-independent, the learning process is domain independent, and the learned heuristic is more effective than the standard FF heuristic.

In this paper, we improved on these results by framing the learning problem as a learning to rank problem instead of an ordinary regression problem. This is motivated by the insight that, in a greedy search, the ranking induced by a heuristic, rather than its numerical values, governs the success of the planning. By optimizing for the ranking directly, our RankSVM learner is able to produce a heuristic that outperforms heuristics learned through least-squares regression.

Additionally, we introduce new methods for constructing features for heuristic learners. Like Yoon et al., we derive our features from an existing domain-independent heuristic. However, our features focus on the ordering and interaction between actions in approximate plans. Thus, they can be based on any existing heuristic that implicitly constructs an approximate plan, such as the context-enhanced additive (CEA) heuristic. These features can be easily constructed and still encode a substantial amount of information for heuristic learners.

In our experiments, we evaluated the performance of the different configurations of our learners on several of the International Planning Competition learning track problems. We find that the learned heuristics using the RankSVM approach allow more problems to be solved successfully than using the popular FF and CEA heuristics alone. Additionally, they significantly surpass the performance of heuristics learned through ordinary regression.

Bounded Optimal Exploration in MDP[7]

Within the framework of probably approximately correct Markov decision processes (PAC-MDP), much theoretical work has focused on methods to attain near optimality after a relatively long period of learning and exploration. However, practical concerns require the attainment of satisfactory behavior within a short period of time.

In this paper, we relaxed the PAC-MDP conditions to reconcile theoretically driven exploration methods and practical needs. We propose simple algorithms for discrete and continuous state spaces, and illustrate

the benefits of our proposed relaxation via theoretical analyses and numerical examples. Our algorithms also maintain anytime error bounds and average loss bounds. Our approach accommodates both Bayesian and non- Bayesian methods.

Optimization as Estimation with Gaussian Processes in Bandit Settings[11]

Recently, there has been rising interest in Bayesian optimization the optimization of an unknown function with assumptions usually expressed by a Gaussian Process (GP) prior. We have studied an optimization strategy that directly uses an estimate of the argmax of the function. This strategy offers both practical and theoretical advantages: no tradeoff parameter needs to be selected, and, moreover, we establish close connections to the popular GP-UCB and GP-PI strategies.

Our approach can be understood as automatically and adaptively trading off exploration and exploitation in GP-UCB and GP-PI. We illustrate the effects of this adaptive tuning via bounds on the regret as well as an extensive empirical evaluation on robotics and vision tasks, demonstrating the robustness of this strategy for a range of performance criteria.

Sample-Based Methods for Factored Task and Motion Planning[5]




Figure 9: (Left) Each object has a specified goal pose. The initial placements are randomly generated. (Right) The goal is that a single blue object be moved to a different table. The blue object starts at the center of the visible table, and the red objects are randomly placed on the table.

Many important robotic domains of interest require planning in a very high-dimensional space that includes not just the robot configuration, but also the configuration of the external world state, including a variety of quantities such as object poses, reaction states of chemical or biological processes, or intentions of other agents. There has been a great deal of progress in developing probabilistically complete sampling-based methods that move beyond motion planning to multi-modal problems including various forms of task planning. These new methods each require a new formulation, definition of robust feasibility, sampling methods, and search algorithm. This paper presents a general-purpose formulation of a large class of discrete-time planning problems, with continuous or hybrid state and action spaces.

The primary theoretical contribution of this paper is a formulation of factored transition systems that exposes the topology of their solution space, particularly in the presence of dimensionality-reducing constraints. The key insight is that, in some cases, the intersection of solution constraint manifolds is itself a manifold that can be identified using only the individual constraint manifolds. By understanding the topology of the solution space, we can define a robust feasibility property that characterizes a large class of problems for which sampling-based planning methods can be successful.

The primary algorithmic contribution is the construction of two sample-based planning algorithms that exploit the factored, compositional structure of the solution space to draw samples from a space in which solutions have positive measure. These algorithms search in a combined space that includes the discrete structure (which high-level operations, such as pick or place happen in which order) and parameters (particular continuous parameters of the actions) of a solution. Theoretically, these algorithms are probabilistically complete when given sufficient samplers. Practically, they can solve complex instances of task-and-motion planning problems (such as shown in Figure 9).

Provably Safe Robot Navigation with Obstacle Uncertainty[1]

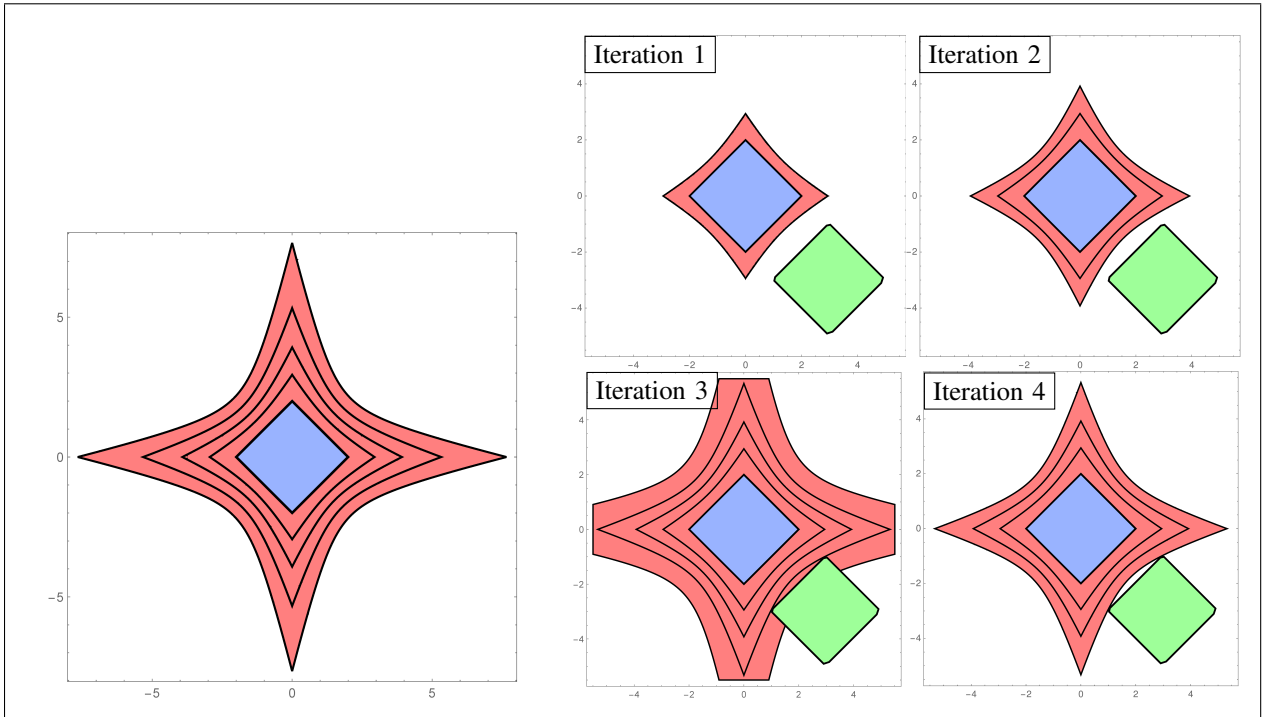


Figure 10: (Left) The blue square represents the mean estimated obstacle. Each outline in the red set represents a different probability shadow of the obstacle. (Right) A line search for the maximal shadow. The shadow grows and shrinks until it contacts the green space visited by the robot.

Safe and reliable operation of a robot in a cluttered environment can be difficult to achieve due to noisy and partial observations of the state of both the world and the robot. As autonomous systems leave the factory floor and become more pervasive in the form of drones and self-driving cars, it is becoming increasingly important to understand how to design systems that will not fail under these real-world conditions. While it is important that these systems be safe, it is also important they do not operate so conservatively as to be ineffective. They must have a strong understanding of when they take risks so they can avoid them, but still operate efficiently.

While most previous work focuses on robot state uncertainty, this paper focuses on safe navigation when the locations and geometries of these obstacles are uncertain. We focus on algorithms that find safety

certificates easily verifiable proofs that the trajectory or policy is safe. We examine two implications of the algorithms. First, the computational complexity of reasoning about uncertainty can be quite low. Second, the mathematics surrounding robot safety can have surprising behavior. We demonstrate how these tools can be used to design a motion planner guaranteed to give only safe plans, and inform the design of more general systems that make decisions under uncertainty.




Figure 11: The desired trajectory found by the planner shown with its specialized shadows that certify the probability of collision as less than 0.26%

This paper makes three contributions. The first is a formal definition of online safety that provides risk bounds on the entire execution of a policy. The second contribution is an algorithm for efficiently verifying offline safety with respect to polytopes with Gaussian distributed faces (PGDFs) that is then generalized to the online case (see Figure 10). In comparison to previous methods, the quality of the resulting bound is not dependent on the number of obstacles in the environment. The presented algorithms produce a certificate, which allows another system to efficiently verify that the actions about to be taken are safe. For a maximal collision probability of ϵ , the runtime of the algorithm grows as $\log \frac{1}{\epsilon}$ making it efficient even for very small ϵ s.

The third contribution is a modification to the RRT algorithm that generates safe plans. For any fixed ϵ , the resulting planner is guaranteed to only return trajectories for which the probability of failure is less than ϵ . We note that for n obstacles, the runtime of the RRT is increased only by a $\log n \log \frac{1}{\epsilon}$ factor, which suggests that reasoning about uncertainty can come at a low computational cost. A result of running this algorithm is shown in Figure 11.

Decidability of Semi-Holonomic Prehensile Task and Motion Planning[2]

The last few decades of robotic planning have been dominated by sample-based techniques. Sample-based techniques are very useful tools to quickly find solutions in many domains. However, they suffer from the notable drawback that they cannot prove that a solution does not exist for a particular problem.

The existence of a probabilistically complete algorithm for a planning problem does not settle the question of whether a complete decision procedure, an algorithm that indicates whether a solution does or does not

exist for any problem instance, exists. For classic motion planning, a holonomic robot among static obstacles, we know that exact algorithms exist for the general case. However, for motion planning in the presence of movable objects, the results are much more limited.

The formal treatment of the problem of planning among movable objects was initiated by Wilfong. When the number of placements and grasps is finite, the problem can be shown to be decidable by building a manipulation graph consisting of a finite number of transfer and transit paths. Decidability for continuous grasps and placements, but involving a single movable object, was shown by Dacre-Wright et al.. More recently, decidability was shown for planning with two objects under restrictive geometries and dynamics.

In this paper, we consider a much more general version of planning in the presence of movable obstacles. We allow an arbitrary dimensional world with an arbitrary number of robots, objects, and obstacles, all with semi-algebraic geometries. We also assume that each robot can be holonomically controlled, and each object can be holonomically manipulated. In this manner, we can account for various continuous polynomial dynamics including translations, rotations, stretching, twisting, and morphing. We do restrict our attention to prehensile manipulation, where objects are rigidly attached to appropriate robots during manipulation. We call the resulting class of problems prehensile task and motion planning (PTAMP).

We define a general task and motion planning framework capable of representing a large variety of planning problems including PTAMP. At the core of the framework is the concept of semi-holonomic controllability (SHC), which accurately describes the intrinsic dynamics of many task and motion planning problems including PTAMP.

The central result of the paper is: jointly-controllably-open (JC-open) domains are decidable. We then give a perturbation algorithm and show that any real-world PTAMP can be rewritten to be JC-open.

We give a constructive proof of the decidability of JC-open domains. Our algorithm is divided in four parts. First, we use a decomposition algorithm that decomposes the configuration space into a finite number of manifolds with special properties. Next, we use techniques from differential geometry to calculate the internal controllability of each manifold. Afterwards, for every manifold, we calculate its stratified controllability, i.e. the controllability gained by leaving a manifold and utilizing the controllability of neighboring manifolds. To accomplish this step, we present the convergence condition, which we shows holds for JC-open domains. Finally, we execute a graph search to calculate the reachability set for our initial configuration and test for the existence of a solution.

References

- [1] Brian Axelrod, Leslie Pack Kaelbling, and Tomas Lozano-Perez. Provably safe robot navigation with obstacle uncertainty. In *Robotics: Science and Systems (RSS)*, 2017.
- [2] Ashwin Deshpande, Leslie Pack Kaelbling, and Tomas Lozano-Perez. Decidability of semi-holonomic prehensile task and motion planning. In *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.
- [3] Caelan Reed Garrett, Leslie Pack Kaelbling, and Tomas Lozano-Perez. Learning to rank for synthesizing planning heuristics. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2016.
- [4] Caelan Reed Garrett, Tomas Lozano-Perez, and Leslie Pack Kaelbling. Backward-forward search for manipulation planning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.
- [5] Caelan Reed Garrett, Tomas Lozano-Perez, and Leslie Pack Kaelbling. Sample-based methods for factored task and motion planning. In *Robotics: Science and Systems (RSS)*, 2017.
- [6] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Implicit belief-space pre-images for hierarchical planning and execution. In *IEEE Conference on Robotics and Automation (ICRA)*, 2016.
- [7] Kenji Kawaguchi. Bounded optimal exploration in MDP. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1758–1764. AAAI Press, 2016.

- [8] G.D. Konidaris, L.P. Kaelbling, and T. Lozano-Perez. Symbol acquisition for probabilistic high-level planning. In *Proceedings of the Twenty Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [9] Gilwoo Lee, Tomas Lozano-Perez, and Leslie Pack Kaelbling. Hierarchical planning for multi-contact non-prehensile manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.
- [10] Xinkun Nie, Lawson L.S. Wong, and Leslie Pack Kaelbling. Searching for physical objects in partially known environments. In *IEEE Conference on Robotics and Automation (ICRA)*, 2016.
- [11] Zi Wang, Bolei Zhou, and Stefanie Jegelka. Optimization as estimation with Gaussian processes in bandit settings. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.
- [12] Lawson L.S. Wong, Thanard Kurutach, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Object-based world modeling in semi-static environments with dependent Dirichlet process mixtures. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.